

Tweeting Behaviour during Train Disruptions within a City

Nut Limsopatham

M-Dyaa Albakour

Craig Macdonald

Iadh Ounis

School of Computing Science
University of Glasgow
firstname.lastname@glasgow.ac.uk

Abstract

In a smart city environment, citizens use social media for communicating and reporting events. Existing work has shown that social media tools, such as Twitter and Facebook, can be used as social sensors to monitor events in real-time as they happen (e.g. riots, natural disasters and sport events). In this paper, we study the reactions of citizens in social media towards train disruptions within a city. Our study using 30 days of tweets in a large city shows that citizens react differently to train disruptions by, for instance, displaying unique behaviours in tweeting depending on the time of the disruption. Specifically, for working days, tweets related to train disruptions are typically generated during rush hour periods. In contrast, during weekends, urban citizens tended to tweet about train disruptions during late evenings. Using these insights, we develop a supervised approach to predict whether a train disruption tweet will be retweeted and propagated on the social network, by using features, such as time, user, and the content of tweets. Our experimental results show that we can effectively predict when a train disruption tweet is retweeted by using such features.

Introduction

In the smart cities of the 21st century, citizens have a variety of emerging information needs in their public urban spaces (Albakour et al. 2014; Kukka et al. 2013). Social media platforms are important elements of the city's knowledge infrastructure that can assist in serving the information needs of the citizens, where they communicate with each other and report on news or event. For example, studies have proposed to use the Twitter social media platform as a source for detecting and retrieving real-world events (McCreadie et al. 2013; Metzler, Cai, and Hovy 2012). Therefore, such platforms have the potential to be valuable data sources to provide information and insights about other elements of a smart urban environment, such as transportation networks, where they can be used as sensors (Sheth 2009).

Indeed, intelligent transportation networks are important components of a smart urban community (An, Lee, and Shin 2011). Reliable and real-time transportation information are

among the emerging information needs of citizens and authorities in smart cities (Kukka et al. 2013). Traditionally, such networks have been managed and optimised with traditional data sources (Ran and Boyce 1994). On the other hand, user-generated content on social media streams can play an important role in providing more context about the operation of transportation networks or otherwise in predicting transportation-related incidents. In particular, many social media users report on transportation disruptions in real-time (Sasaki et al. 2012) and there have been some studies on transportation disruptions on social media. For example, Twitter has been used to provide more contextual information about traffic incidents on motorways and major road networks (Daly, Lecue, and Bicer 2013). Sasaki et al. (2012) proposed to use a simple thresholding over the volume of tweets, which are related to train disruptions, in order to detect the status of the train service in Japan.

In this paper, we also study the use of Twitter in the context of train disruption. In contrast to predicting disruptions from the tweets of the users, we aim to uncover the temporal patterns of tweeting about train disruptions. While such patterns may be of interest to social scientists studying citizens' behaviours during travel disruption, in this paper we show that they can also assist in a tweet classification task.

Furthermore, we study how train disruption tweets propagate, by being retweeted, in the Twitter social network. Retweeting is a method of sharing important or interesting information with the user's community (Petrovic, Osborne, and Lavrenko 2011). The understanding of retweeting behaviours would uncover what makes the user's community interested in a particular piece of information (Naveed et al. 2011). Indeed, there have been many studies on analysing and predicting tweet propagation on the Twitter social network. Luo et al. (2013) studied the problem of predicting who is going to retweet a user. In their study, the retweet history of the user and the content of the tweet are among the best performing features. Petrovic et al. (2011) performed a user study to predict which tweets are likely to be retweeted. They proposed a supervised approach that performed as well as humans in predicting retweets and reported that social features about the user are among the most effective for this task. Naveed et al. (2011) also used a supervised approach and conducted a study of what makes a tweet worth retweeting. They found that general topics and bad news are more

likely to be retweeted than personal issues. Similarly, we aim to predict whether a train disruption tweet will be retweeted. These tweets represent the interesting information about train disruptions that users share and would give insight into the topics that are of interest to the users when it comes to train disruptions. For example, such topics may include the durations of delays of the trains or whether there is a replacement bus. In this paper, we develop a supervised machine learning approach, using general and disruption-related features about the content of the tweets, the users and the time to predict whether a train disruption tweet is retweeted.

Using 30 days of tweets that are collected within a large urban city, we conduct a study of the temporal tweeting behaviour about train disruptions. We show that train disruption tweets exhibit unique temporal daily patterns as opposed to the general tweeting behaviour. In addition, such patterns differ between weekdays and weekends. We also use this dataset to thoroughly evaluate our supervised retweet prediction approach, and to identify the important features for retweet prediction in a train disruption context. We show that state-of-the-art features for retweet prediction, such as the user credibility features, do not necessarily work in a train disruption context. Instead, disruption-related features and the content of the tweets are among the most effective for this task. These features surface the important topics that citizens may be particularly interested in when it comes to train disruptions, for example, the mention of the time of delayed trains.

Datasets

We study the patterns of tweets and retweets about train disruption within Glasgow, a large city in the United Kingdom. We collected tweets generated between 24 November 2014 and 23 December 2014 (30 days) by using the Twitter Streaming API¹. In particular, using the Streaming API, we filter only tweets related to the Glasgow city (as demonstrated by the geo-locations of tweets and the mentioning of keywords, such as Glasgow).

In order to uncover unique patterns of tweets about train disruption, we generated two datasets from this Twitter stream. The first dataset, called the GLASGOW dataset, contains all of the aforementioned Glasgow tweets. Next, from the GLASGOW dataset, we followed Sasaki et al. (2012) and identified train disruption tweets by filtering tweets related to both train (e.g. mentioning one of the sixty two train stations in Glasgow) and disruption (e.g. mentioning terms, such as “delay”, “late” and “cancel”). We refer to this second dataset as the DISRUPTION dataset. We manually validated 100 tweets sampled from the DISRUPTION dataset and found that 83% were related to train disruption.

Table 1 shows the statistics of tweets from both the GLASGOW and DISRUPTION datasets. The GLASGOW dataset contains 3.58 million tweets. On average, there were 119.2k tweets related to the Glasgow city generated each day. On the other hand, the DISRUPTION dataset, which is a subset of the GLASGOW dataset, has 1.7k tweets generated during the same period. The number of tweets in the DISRUPTION

Dataset	#tweets in total	Avg. #tweets per day	Min. #tweets per day	Max. #tweets per day
GLASGOW	3.58M	119.2k	99.1k	166.7k
DISRUPTION	1.7k	56.83	8	192

Table 1: The statistics of tweets from the datasets.

dataset ranges between 8 and 192 tweets per day, with an average of 56.83 tweets.

Analysing Patterns of Train Disruption Tweets


In order to gain a better understanding of the use of tweets in reporting train disruptions within Glasgow, we compare train disruption tweets (i.e. tweets from the DISRUPTION dataset) with general tweets (i.e. tweets from the GLASGOW dataset), all of which were generated during the same period. In particular, we aggregate the volume of tweets from each dataset and calculate the distribution of the likelihood that a tweet from each dataset is generated during each hour of the day.

Figure 1(a) shows the distribution of tweets from both the GLASGOW and the DISRUPTION datasets. We observe that tweets from the two datasets have different patterns. In particular, tweets related to the Glasgow city (i.e. from the GLASGOW dataset) are more likely to be generated during the day time than at night. Meanwhile, the train disruption tweets have a unique pattern, in that they are more likely to be generated during ‘rush hour’ periods, especially between 06:00 and 09:00. This is correlated with the time that most people commute to work. Next, in Figures 1(b) and 1(c), we compare the tweet patterns in weekdays and weekends for the GLASGOW and the DISRUPTION datasets, respectively. From Figure 1(b), for the GLASGOW dataset, we observe that the patterns of tweets generated during weekdays and weekends are very similar. In contrast, as shown in Figure 1(c), the likelihood distributions that train disruption tweets were generated in weekdays and weekends were different. In particular, during weekdays, tweets about train disruption were likely to be generated during the rush hours, while during weekends most of the train disruption tweets are generated in late evening (e.g. between 17:00 to 21:00). These may reflect the behaviour of train users that take a train to work in weekdays and go out at late night in weekends. These observations suggest that the pattern of tweets about train disruption was unique and, in general, predictable. In the next section, we study the propagation of train disruption messages via retweeting.


Predicting Train Disruption Retweets

In this section, we study the propagation of train disruption tweets. We aim to automatically predict a retweet of a train disruption tweet (i.e. *retweetability*). We consider this problem as a classification task. We use a linear SVM classifier (Chang and Lin 2011). In particular, we follow Petrovic et al. (2011) and train a binary classifier with positive and negative examples of train disruption tweets, as to whether or not they were retweeted. To build the classifier, we investigate three main sets of features that can be directly ex-


¹<https://dev.twitter.com/streaming/public>



(a) The likelihood distribution of tweets generated each hour in the GLASGOW and the DISRUPTION datasets.



(b) The likelihood distribution of tweets generated in weekdays and weekends from the GLASGOW datasets.



(c) The likelihood distribution of tweets generated in weekdays and weekends from the DISRUPTION datasets.

Figure 1: The likelihood distribution of tweets generated at different hours of the day.

tracted from each tweet without significant computing effort. We do not extract features related to the individual terms used in the tweets or the user’s profile. Specifically, the three sets of features consist of *user features* (9), *tweet features* (8) and *time-related features* (6).

User features. The first set of features are extracted from the user information, including (1) *whether or not the user was verified*, (2) *whether or not the user provides a URL*, (3-6) *number of followers, friends, favourites, statuses*, (7) *number of times the user was listed*, (8) *number of days the user has been a member of Twitter* and (9) *number of characters in the user description*. This set of features focus on the credibility of the user. Petrovic et al. (2011) found that 91% of tweets from verified users were retweeted; hence, *whether or not the user was verified* is likely to be a good feature. Suh et al. (2010) showed that the number of followers and friends were good indicators of retweetability, while the number of favourites and statuses were weaker features. A list is a way to follow a user. If a user is listed, the tweets generated by this user will be directed to other users who have the user in their list. Hence, if the user is listed by many

other users, tweets from this user are likely to be seen and retweeted. We also explore the effect of providing a URL, a description and the number of days the user has been a member of Twitter on retweetability. Providing a URL and description in the user profile is one of the ways to improve the credibility of the user, as it allows other users to explore more information about this particular user. The number of days the user has been a member of Twitter may indicate the level of experience with Twitter.

Tweet features. The second set of features are related to the tweet itself, which include (1) *the length of the tweet (i.e. number of characters)*, (2) *whether or not the tweet is in English*, (3) *whether or not the tweet is a reply*, (4) *whether or not the tweet is a retweet*, (5-8) *the number of hashtags, mentions, URLs and trending words*. Very short tweets may not be informative; hence, they are unlikely to be retweeted. Petrovic et al. (2011) showed that tweets written in English were more likely to be retweeted, while a replied or retweeted tweet was less likely to be retweeted. Trending words are words in the tweet that Twitter considers to be trending at that particular time, so if the tweet contains trending words, it is likely to be retweeted. The number of hashtags, mentions and URLs used in the tweet is highly correlated with retweetability (Suh et al. 2010).

Time-related features. The third set of features are inspired by the unique pattern of train disruption tweets. As the tweet pattern is different across the time of the day and the day of the week that the tweet is generated, we extract the following time-related features from each tweet: (1) *the hour of day* and (2) *the day of week when the tweet is created*, (3-4) *is the tweet created on weekdays or weekends*, (5) *whether or not the tweet content contains a time (e.g. “14:20”)* and (6) *whether or not the tweet content contains a number*. From the analysis of patterns of train disruption tweets, we found that users were likely to tweet a lot during specific hours; hence, it was more likely that train disruption tweets generated during those hours would be retweeted (see Figure 1(a)). In addition, we found that the tweet pattern of a train disruption was different based on the day of the week (i.e. weekdays vs. weekends). Therefore, it is intuitive to use this information as features for retweet prediction. In addition, commuting users typically focus on how many minutes the train will be delayed or what is the new time schedule of the train; hence, whether the tweet content contains numbers or time could be effective features.

Experiments

To evaluate our approach to automatically predict train disruption retweets, we use the previously described DISRUPTION dataset. We divide the DISRUPTION dataset into a *training* and a *testing* subset. The training subset consists of tweets from the first twenty three days, while tweets from the last seven days are used for testing. Table 2 shows the total number of tweets and the number of tweets that were retweeted from the dataset. Over the three datasets, 44% to 52% of the tweets were retweeted.

We compare the performance of our approach for predicting train disruption retweets with both a random and a majority baseline. The random baseline randomly decides

Dataset	#retweeted	#not retweeted	#total
DISRUPTION	851 (50%)	854 (50%)	1,705
DISRUPTION (training)	642 (52%)	592 (48%)	1,234
DISRUPTION (testing)	209 (44%)	262 (56%)	471

Table 2: Number of retweeted tweets in the datasets.

Model	F_1	Accuracy	Precision	Recall
Random	0.4934	0.5074	0.4538	0.5407
Majority	0.6147	0.4437	0.4437	1.0
SVM (All)	0.7357	0.7346	0.6591	0.8325
SVM (All\User)	0.9499	0.9554	0.9522	0.9476
SVM (All\Tweet)	0.4348	0.5308	0.4067	0.4670
SVM (All\Time)	0.4030	0.4968	0.4255	0.3827

Table 3: Performance of retweet prediction.

whether a tweet will be retweeted, while the majority baseline classifies all tweets as positive (i.e. will be retweeted) as the majority of tweets in the training set are retweeted. Table 3 compares the performance of our approach (*SVM (All)*) with the two baselines, in terms of F_1 and Accuracy. We observe that our approach markedly outperforms both the random and majority baselines, in terms of both F_1 and Accuracy. Indeed, our approach attains an F_1 score of 0.7357.

Next, we analyse the effect of each set of features used. We do an ablation study by removing one set of features out when learning a classifier. For example, *SVM (All\User)* is when we use only the tweet and time-related features without the user features. We observe that the user features harm the performance of our classifier. In particular, without using the user features, which were effective features in the literature (e.g. (Petrovic, Osborne, and Lavrenko 2011; Suh et al. 2010)), the F_1 score further improves to 0.9499. Meanwhile, the tweet and time-related features are effective. As we remove each of the set of tweet and time-related features, the F_1 score is reduced to 0.4348 and 0.4030, respectively. The poor performance of the user features is in contrast to existing work on retweet prediction (Petrovic, Osborne, and Lavrenko 2011; Suh et al. 2010), where it was shown that features extracted from the user profile are typically effective. This shows a different retweeting behaviour in the case of train disruption. Indeed, users retweeting train disruption focus more on the time and content of the tweets rather than who did post the tweet.

In addition, we test our approach by predicting retweets from the generic GLASGOW dataset, by using the same experimental settings. We found that our approach does not perform very well ($F_1 = 0.2945$). This is expected as we use only a subset of features from the existing approaches and use some new features that are designed specifically to train disruption tweets. However, our overall conclusion is that retweets about train disruption have a unique characteristic that can be effectively modelled by using the content and time extracted directly from the tweets themselves.

Conclusions

In this work, we studied the patterns of train disruption tweets in Glasgow. We observed unique tweeting and retweeting behaviours about train disruptions. From the analysis of the likelihood of tweeting about train disruptions

over different hours of a day, in contrast to tweets in general, we found that users tweeted about train disruptions mostly during rush hour periods. Furthermore, using insights from this analysis, we developed disruption-related features for training a classifier to predict train disruption retweets. We showed that our approach was effective, as it achieved an F_1 score of 0.9499. From the analysis of the results, we also found that the retweetability of train disruption tweets were different from tweets in general. In particular, we found that, in contrast to existing work that predicted generic retweets, user features were not useful for predicting train disruption retweets. On the other hand, we found that features about the time and content of the tweet were effective for this task.

Acknowledgements

The authors acknowledge the support of the Economic and Social Research Council who funded the integrated Multi-media City Data project, on which this paper is based.

References

- Albakour, M.; Macdonald, C.; Ounis, I.; Clarke, C. L. A.; and Bicer, V. 2014. Information Access in Smart Cities (i-ASC). In *ECIR'2014*.
- An, S.-h.; Lee, B.-H.; and Shin, D.-R. 2011. A Survey of Intelligent Transportation Systems. In *CICSyN'2011*.
- Chang, C.-C., and Lin, C.-J. 2011. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology* 2. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Daly, E. M.; Lecue, F.; and Bicer, V. 2013. Westland Row Why So Slow? Fusing Social Media and Linked Data Sources for Understanding Real-Time Traffic Conditions. In *IUI'2013*.
- Kukka, H.; Kostakos, V.; Ojala, T.; Ylipulli, J.; Suopajarvi, T.; Jurmu, M.; and Hosio, S. 2013. This is not classified: everyday information seeking and encountering in smart urban spaces. *Personal and Ubiquitous Computing* 17(1).
- Luo, Z.; Osborne, M.; Tang, J.; and Wang, T. 2013. Who Will Retweet me? Finding Retweeters in Twitter. In *SIGIR'2013*.
- McCreadie, R.; Macdonald, C.; Ounis, I.; Osborne, M.; and Petrovic, S. 2013. Scalable Distributed Event Detection for Twitter. In *IEEE Big Data'2013*.
- Metzler, D.; Cai, C.; and Hovy, E. 2012. Structured Event Retrieval over Microblog Archives. In *NAACL HLT'2012*.
- Naveed, N.; Gottron, T.; Kunegis, J.; and Alhadi, A. C. 2011. Bad News Travel Fast: A Content-based Analysis of Interestingness on Twitter. In *WebSci'2011*.
- Petrovic, S.; Osborne, M.; and Lavrenko, V. 2011. RT to Win! Predicting Message Propagation in Twitter. In *ICWSM'2011*.
- Ran, B., and Boyce, D. 1994. *Dynamic Urban Transportation Network Models: Theory and Implications for Intelligent Vehicle-Highway Systems*. Number 417.
- Sasaki, K.; Nagano, S.; Ueno, K.; and Cho, K. 2012. Feasibility Study on Detection of Transportation Information Exploiting Twitter as a Sensor. In *ICWSM'2012*.
- Sheth, A. 2009. Citizen Sensing, Social Signals, and Enriching Human Experience. *IEEE Internet Computing* 13(4).
- Suh, B.; Hong, L.; Pirolli, P.; and Chi, E. H. 2010. Want to Be Retweeted? Large Scale Analytics on Factors Impacting Retweet in Twitter Network. In *SOCIALCOM'2010*.